• V-Dem

...

.

...........

.....

.

.........

.

ė

......

.

......

.........

....

•••••

.....

000

••

.

.....

....

...

.................

.....

....

.

......

••

.....

On Revolutions

......

...

.

Armand M. Leroi, Ben Lambert, Matthias Mauch, Marina Papadopoulou, Sophia Ananiadou, Staffan I. Lindberg, Patrik Lindenfors

. ė •• 0000 ā ē ۲ •• ė ò ò ò õ ē • ē •••

.

.....

.....

....

...

...

March 2018

Working Paper SERIES 2018:63

THE VARIETIES OF DEMOCRACY INSTITUTE



UNIVERSITY OF GOTHENBURG dept of political science Varieties of Democracy (V–Dem) is a new approach to conceptualization and measurement of democracy. It is co-hosted by the University of Gothenburg and University of Notre Dame. With a V–Dem Institute at University of Gothenburg with almost ten staff, and a project team across the world with four Principal Investigators, fifteen Project Managers (PMs), 30+ Regional Managers, 170 Country Coordinators, Research Assistants, and 2,500 Country Experts, the V–Dem project is one of the largest ever social science research-oriented data collection programs.

Please address comments and/or queries for information to:

V-Dem Institute Department of Political Science University of Gothenburg Sprängkullsgatan 19, PO Box 711 SE 40530 Gothenburg Sweden E-mail: contact@v-dem.net

V-Dem Working Papers are available in electronic format at www.v-dem.net. Copyright ©2018 by authors. All rights reserved.

On Revolutions^{*}

Armand M. Leroi[†], Ben Lambert[‡], Matthias Mauch[§], Marina Papadopoulou,[¶] Sophia Ananiadou^{||}, Staffan I. Lindberg^{*}, Patrik Lindenfors^{††}

^{*}B.L. was supported by EPSRC grant code: EP/F500394/1. V-Dem data collection was supported by European Research Council, Grant 724191; Riksbankens Jubileumsfond, Grant M13-0559:1; Knut and Alice Wallenberg Foundation to Wallenberg Academy Fellow, Grant 2013.0166; as well as by internal grants from the Vice- Chancellor's office, the Dean of the College of Social Sciences, and the Department of Political Science at University of Gothenburg. P.L. was also supported by the Marianne and Marcus Wallenbergs Foundation Grant 2017.0049

[†]Corresponding author: a.leroi@imperial.ac.uk; Department of Life Sciences, Imperial College London, London SW7 2AZ, UK, & Data Science Institute, Imperial College London, London SW7 2AZ, UK

[‡]Department of Zoology, University of Oxford, New Radcliffe House, Woodstock Road, Oxford OX2 6GG, UK & Department of Infectious Disease Epidemiology, Imperial College London, Medical School Building St Mary's Campus, Norfolk Place, London W2 1PG, UK

[§]Centre for Digital Musicology, Queen Mary University, London, Mile End Road, London E1 4NS, UK [¶]Department of Life Sciences, Imperial College London, London SW7 2AZ, UK, & Groningen Institute

for Evolutionary Life Sciences, University of Groningen, Nijenborgh 7, Groningen 9747 AG, Netherlands

^{||}National Centre for Text Mining, School of Computer Science, University of Manchester, 131 Princess Street, Manchester, M1 7DN, UK

^{**}V-Dem Institute, Gothenburg University, Box 100, Gothenburg SE-405 30, Sweden & Department of Political Science, Gothenburg University, Box 100, Gothenburg S-405 30, Sweden

^{††}Institute for Future Studies, Box 591, SE-101 31 Stockholm, Sweden; Centre for the Study of Cultural Evolution & Department of Zoology, Stockholm University, SE-106 91 Stockholm, Sweden

Abstract

Sometimes the normal course of events is disrupted by a particularly swift and profound change. Historians have often referred to such changes as "revolutions" and, though they have identified many of them, they have rarely supported their claims with statistical evidence. Here we present a method to identify revolutions based on a measure of the multivariate rate of change called Foote Novelty. We define revolutions as those periods of time when the value of this measure, F, can, by a non-parametric test, be shown to be significantly greater than the background rate. Our method also identifies conservative periods when the rate of change is unusually low. Importantly, our method permits searching for revolutions over any time scale that the data permit. We apply it to several quantitative data sets that capture long-term political, social and cultural changes and, in some of them, identify revolutions, both well known and not. Our method is a general one that can be applied to any phenomenon captured by multivariate time series data of sufficient quality.

What is a revolution?

It seems that the word "revolution" was first applied to sublunary events when parliamentarians, aided by the Dutch, deposed James II from the English throne and so brought about the Glorious Revolution. Since then, it has been applied ever more widely (Cohen 1986). Responding to the French Revolution of 1789, Friederich Schlegel called for an Aesthetic Revolution in poetry, and so extended the term beyond politics (Heumakers 2015). In the latter half of the 19th century — an age of revolutions — John Stuart Mill, Karl Marx and Arnold Toynbee, following a French coinage, wrote of the Industrial Revolution (Bezanson 1922). In the 1950s Alexandre Koyré (1957), Herbert Butterfield (1950), A. R. Hall (1954) and Thomas Kuhn (1957), descried the Scientific Revolution (Cohen 1994). In The Structure of Scientific Revolutions, Kuhn (1972) generalized the idea, arguing that science advanced, if it advanced at all, by revolutions. The Darwinian Revolution was swiftly identified (Ruse 1979; Himmelfarb 1996), as were many others. Indeed, Kuhn's book prompted something of a revolution in scientific discourse, as scientists themselves took to identifying, or calling for, "paradigm shifts" — Kuhn's term for a revolution — in their fields. A search of all articles indexed by the Web of Science in 2017 reveals more than two thousand that do so, though many of the purported revolutions seem quite modest in scope (e.g., Seward 2017; Raoult 2017; Lowenstein and Grantham 2017; Lonne 2017).

For all that, revolutions are hard to pin down. Upon close inspection they often seem to shrink. Pick a revolution, even a famous and well-documented one, and you can be sure to find scholars who have sought to cut it down to size or even deny that it happened at all. "The drastic social changes imputed to the Revolution, seem less clear-cut or not apparent at all." — thus Simon Schama (1989) on how his generation of historians viewed the impact of the French Revolution. "There was no such thing as the Scientific Revolution, and this is a book about it." — so Steve Shapin (1996), in paradoxical mode, on early modern science. Evolutionary biologists may be surprised to learn that the Darwinian Revolution has its skeptics too (Hodge 2005; Bowler 1988).

The difficulty of identifying revolutions has plagued the historical natural sciences as well. In the 1980s archaeologists labelled the sudden appearance, fifty thousand years ago, of culture as the Human Revolution (Mellars and Stringer 1989). It wasn't long before others had dismissed it as the "revolution that wasn't" (Mcbrearty and Brooks 2000). For much of his life Stephen Jay Gould (2002) argued that the Darwinian Revolution had run its course and that evolutionary biology needed another. (But one not to be confused with the broader Paleobiological Revolution of the 1970s and 80s which he helped shape (Sepkoski 2012; Sepkoski and Ruse 2009).) The coping stone of Gould's new paradigm, an unstable edifice, was the theory of punctuated equilibrium that he proposed with Niles Eldredge (Eldredge and Gould 1972). This theory, shorn of its theoretical structure, postulated that change in fossil lineages is itself best described as a series of revolutions rather than gradual evolution. It may seem like a simple matter to decide which, but the ensuing decades-long quarrel among palaeontologists about what the fossils show has proved otherwise (Pennel et al. 2013). Unsurprisingly given its fame, the term "punctuated equilibrium" has flown free from biology and now appears in fields as remote from palaeontology as management science and policy research (e.g., Flink 2017; Fowler et al. 2017; De Ruiter and Schalk 2017). In these fields the term has lost its deeper meaning altogether and is just another way to express the existence, or hope, of revolutions.

The problem is clear. Great revolutions may entail change in many dimensions — ideas, wealth, social roles, political structures, the composition of assemblages of artefacts and species or else their features — but to varying degrees, at varying rates, and with varying starts and ends. A revolution's visibility, then, depends on where you look. Even when considering the same data, some scholars will see discontinuity where others see continuity — it may be merely a matter of temperament — in the absence of an objective method for distinguishing the two, there is no way to know which of their accounts is more true.

In this paper we give the idea of a revolution a statistical foundation. We take the view that, for some set of characteristics of a population that change over time, a revolution is simply a statistically significant local increase in their rate of change relative to the background rate. Given this, our method depends on identifying correlated changes in multivariate time series by means of a non-parametric permutation test. Our work is related to previous work on multivariate time series segmentation (Omranian et al. 2015; Preuss et al. 2015), but differs from these in that it uses the self-similarity of individual series across neighbouring periods (i.e. the local rate of change) to classify time points into "revolutionary" or "conservative" periods, rather than segmenting series into self-contained windows where the series are assumed to follow a time-invariant (typically parametric) relationship. We have previously introduced our method while applying to the evolution of American popular music (Mauch et al. 2015). Here we refine its statistical basis, apply it to several large data sets that capture changes in political, social and cultural systems over time, and identify a variety of revolutions that are well known, as well as some that are not.

Detecting revolutions

A method for detecting revolutions should consider many characteristics of a population simultaneously, that is, be underpinned by a multivariate metric of change. One such method, used in signal processing, is based on a measure called Foote Novelty (Foote 2000), which is based on the collection of pairwise distances between values of the series at different temporal separations, known as a *distance matrix*.

To understand the rationale behind Foote Novelty, consider the one-dimensional time series (1, 2, 2, 1, 5, 4, 5, 4, 4) and its distance matrix (Figure 1). It has an obvious change point after the fourth element, with data before this change (1, 2, 2, 1) relatively homogeneous, and the data afterwards (5, 4, 5, 4, 4) also exhibiting minimal variability. If a distance matrix is calculated from the original data, these two periods manifest as distinct blocks of low local variation along the main diagonal (the darker shaded blocks in Figure 1). By contrast, the pairwise distances between data points before and after such a change are considerable, resulting in two off-diagonal blocks of high-cross variability (the lighter shaded blocks in 1) in the distance matrix. The point of intersection of all four of these blocks occurs at the point of change and, it was Foote's insight that such junctions have the appearance of a checkerboard.

In order to assign high values to time periods that appear checkboard-like and low values to others, Foote devised a kernel that qualitatively resembles a checkerboard. This kernel is composed of two pairs of blocks of size k — the half-width — with the diagonal blocks equal to -1 and the off-diagonal components equal to +1. For example, the kernel for k = 2is given by,

This kernel is moved along the main diagonal and the elementwise product taken of the matrix and the values it overlaps (Figure 1) to calculate a Foote Novelty score, which we argue meters the "revolutionary potential" of a given point in time. Foote Novelty at a time t can, hence, be represented mathematically by the following summation,

$$F_t^k = \sum_{i=-k}^{k-1} \sum_{j=-k}^{k-1} C_{ij}^k D_{t+i,t+j},$$

where C_{ij}^k is the elementwise representation of the Foote kernel C^k with half-width k and D_{ij} represents the elements of the temporally-ordered distance matrix D.

The off-diagonal blocks of C^k contribute positively to this score because they indicate differences in the series before and after the time point under consideration (that in the middle of the kernel) and, hence, can be thought of as calculating the series' "cross-variability" at that time. Since we have no a priori knowledge of the magnitude of revolutions, changes in a series over a period of time are evidence for a revolution only in comparison with a base rate of change. A revolution is characterised by periods of relative stasis punctuated by a period when significant change occurs. If considerable change occurs in either the period immediately preceding the proposed revolutionary epoch or that following it then — at the time scale indicated by the kernel half width — a revolution has not occurred. The diagonal component's of Foote's kernel capture this aspect of a revolution's definition by allowing any local variation in the pre- or post-revolutionary period to diminish the magnitude of the revolution.

For an example of Foote Novelty calculation for our synthetic one-dimensional time series see Figure 1. Note that indices for which the kernel C does not completely overlap with the distance matrix D are omitted. In practice, we use a kernel with two minor modifications. First, we follow Foote in imposing a Gaussian taper with a standard deviation of 0.4k, to remove edge effects. This gives distances closer to the target time point more weight than those further away. Second, in order to have a central point of reference, we add a "cross" of zeros between the blocks of the kernel C. As a result, the size of the whole kernel is 2k + 1, and the value F_t^k corresponds precisely to the kernel centered at t. Assuming that D is large (e.g., 100×100 time points) relative to C (e.g., 4×4), varying t amounts to sliding kernel C along the central diagonal of distance matrix D, calculating F as we go.

In real data, F_t^k is generally positive and varies as the underlying variables fluctuate in value. We therefore define revolutions as periods when its value is statistically significantly higher than in the rest of the series. To determine this we compare the observed F_t^k values to the distribution of F_t^k values obtained from randomly permuting the distance matrix. In our original test we permuted the distance matrix on its axes (Mauch et al. 2015); here, following a suggestion by T. Underwood, H. Long, R. J. So, and Y. Zhu (pers. comm.), we permute on the diagonals.

The kernel half-width, k, can be as small as 1 or as large as the data allow, but different k show different aspects of change. Foote Novelty acts rather like a microscope. Small k values zoom in on short-term heterogeneities that large k values may obliterate, and large k values may reveal long-term variation invisible at smaller fields of view. A sustained period of rapid change will tend to produce revolutionary signals at many different k, but more



Figure 1: Example time series with distance matrix. The 4×4 matrix marked with a blue box shows the natural checkerboard structure of distances at points of rapid change. The corresponding Foote Novelty value F_4^2 is also coloured blue. The offset, present in the original Foote formulation, between the center of the box and the corresponding coordinates of the novelty function is corrected for by our centered formulation.

complex patterns of rate variation will result in conflicting signals. For example, a large k may well identify a single, long, revolution where a smaller one identifies two or more. The shifting picture of the rate landscape that emerges as we adjust the focus of our Foote Novelty microscope is not a weakness of the method, but a consequence of making the scale of analysis explicit. In practice we examine all half-widths that the data allow and identify revolutions by their consistency in a given region.

Any significance value is, of course, arbitrary and we would also like a general picture of fluctuations in the rate of change regardless of whether or not they are statistically significant. To this end we propose an index, R_t , which captures the relative rate of change at a given time point, t. Assuming a set K of desired kernel half-widths, this index is constructed by first standardizing every F_t^k estimate by the average over all valid time points for its halfwidth, \bar{F}^k , and then averaging the standardized values over all $k \in K$ estimated for that time point to give a single value:

$$R_t = \frac{1}{|K|} \sum_{k \in K} \frac{F_t^k}{\bar{F}^k}$$

If $R_t > 1$ then the rate of change at given time point is greater than the average rate of change in the entire series; if $R_t < 1$ smaller.

Having identified a revolution, we'd also like to know which variables contribute to it. One simple way to find out involves removing variables from the data set one at time and re-running the analysis. Variables which, when removed, yield fewer statistically significant F_t^k in a given revolution contribute to it; those which yield more obscure it. Joint effects can be tested by removing combinations of variables. Figure 2A shows the method in action on simulated data. We simulated twenty stationary series, each of which represents a measured variable, for 100 time points (Figure 2A first row). Starting at time point 40 we introduced a revolution by allowing the variables to undergo a directional change for ten generations after which they became stationary again. The rate discontinuity can be clearly seen in the distance matrix (Figure 2A second row). To identify the revolution we estimated F_t^k for all kernel half-widths, k, and time points, t, allowed by the data, in this case $1 \le k \le 49$, and calculated the rate index, R_t , for all time points. A sharp rate discontinuity is visible between time points 37 and 53 where $R_t > 1$ (Figure 2A third row). Finally, we determine the statistical significance for each F_t^k estimate (Figure 2A fourth row). Considering all k, there is strong evidence for a revolution spanning time points 37–57. Since the test's resolution decreases as k increases, the most accurate estimate is given by the smallest k at which the revolution appears, in this case k = 4 where it spans time points 42–49 — very close to the real values of 40–50. A few statistically significant F_t^k values are seen well outside of the simulated revolution; these are false positives and we discuss their identification below.

In this example we simulated stationary variables with a revolution that was both quite strong and long. To see whether our method works in other kinds of series we applied it to several sets of simulated time series and then counted the revolutions detected. In these simulations we varied three parameters set by nature: (1) the *persistence* of the series, ρ , (2) the magnitude of change in variable values during a revolution, that is, its *strength*, *s*, and (3) the *length* of the revolution, *l*. We also varied two parameters set by investigators: (4) the number of variables measured, *n*, and (5) the kernel half-width, *k*. (See Materials and Methods for details). For each combination of parameters, 19,250 in all, we simulated ten replicate populations, and then estimated the rate of false positives (Type I errors) and false negatives (Type II errors).

We investigated the rate of false positives in series with no revolutions (s = 0). In this subset of the simulations only three parameters vary: the persistence of the series, ρ , the number of variables, n, and the kernal half-width, k. Here the overall number of (false) revolutions detected should be equal to, or less than, $\alpha = 0.05$. For fully stationary series $(\rho = 0)$, we found that this was so, however, as the series became more persistent the rate of false positives increased, so that in random walks $(\rho = 1)$, revolutions were detected, on average, in 16% of the series (Figure 3A). Thus autocorrelation produces a large number of false-positives or, to put it another way, like many econometric tests, ours requires stationary series. Persistent series can be made stationary by taking their differences and, when we do so, we find that the rate of false positives is, once again, equal to or below the set significance threshold regardless of persistence (Figure 3A). Figure 2B illustrates the effect of differencing



Figure 2: Identifying revolutions using Foote Novelty in simulated series. A. Evolution of 20 simulated stationary time series with a revolution in the middle. B. Evolution of 20 simulated undirected random walk time series with a revolution in the middle. In both sets of simulations the standard deviation of non-revolutionary periods is set at $\sigma = 1$. During the revolution ends. The amount by which each variable, *i*, changes during during a revolution is drawn from a normal distribution, $d_i \sim \mathcal{N}(0, s)$, where s = 1 is the "revolution strength". First row from top: Evolution of the time series. Second row: Distance matrices among time points: dark blues are increasingly dissimilar. Third row: The rate of change index, R_i , which is the sum of the F_i^k values for any time point *i* over all *k*, relative to the sum of the mean F_i^k values over all time points. Fourth row: Identifying revolutions by Foote Novelty. Each cell represents the F_i^k estimate for a given half-width, *k* and time point; the colour of the cell gives the relative F_i^k value, light grey being low and dark grey being high. Note that this colour scale is only comparable within any given plot. Statistically significant ($\alpha = 0.05/2$) revolutionary periods are overlain in red, conservative periods are blue. In both cases we correctly identify a revolution in the correct region, but at larger half-widths the resolution becomes coarser. Statistically significant time points which are not contiguous with the simulated revolution are false positives. Note that, for the random walk series, the distance matrices, R_i and F_i^k values are all based on first differences. This means that only revolution boundaries are expected to have high F_i^k values.

on one set of random walk time series with a revolution introduced between time points 40– 50. Now the revolution appears as spikes in F_t^k and R_t marking its start and end and a set of significant F_t^k values between time points 32–41 and 46–57. Smaller k values (e.g, k = 8) give the most accurate estimates of the revolution's boundaries as time points 39–41 and 49–50 (Figure 2B third and fourth rows).

We investigated the rate of false negatives in all series which contained revolutions (s > 0). When applied to levels we found that, regardless of persistence, our test fails to detect about 22% of the revolutions (Figure 3B). Differencing reduces the power of the test considerably when applied to stationary series, but only slightly in highly persistent series (Figure 3B). Focusing on the two extreme cases, stationary series ($\rho = 0$) and random walk series ($\rho = 1$) made stationary by differencing we find that our method tends to fail to identify short and



Figure 3: **Performance of revolution detection by Foote Novelty based on simulated series. A.** Mean Type I error (false positive) rate as function of persistence, ρ , in series without revolutions (s = 0). When estimated on levels (solid grey line), the observed Type I error rate quickly increases above the significance level, $\alpha = 0.05$ (solid red line), but when differenced it does not (dotted grey line). **B.** Mean Type II error (false negative) rate as function of persistence in series with revolutions (s > 0). When estimated on levels (solid line), the observed Type II error rate is around 22%, but is higher when estimated on differences, decreasing as persistence increases. **C.** A closer look at Type II error rates in stationary series ($\rho = 0$) (Top) and first-differenced random walk series ($\rho = 1$) (Bottom) as as a function of the kernal half-width, k, number of variables in the simulation, n, the strength of the revolution, s, and its length, l. These plots are an expansion of the data in **B** marked on few variables ($n \le 10$), particularly when analysed using very small half-widths (k = 1). Although the distribution of false negatives differs somewhat between the two sets of series, the overall mean false negative rates are very similar, 22 and 23% respectively.

weak revolution ($l \leq 6$, $s \leq 0.5$), in data sets based on few variables ($n \leq 10$), particularly when analysed using very small half-widths (k = 1) (Figure 3C). In order to balance the risk of Type I and II errors when applying our test to real data we therefore recommend that investigators first estimate the overall persistence, $\hat{\rho}$, of the set of time series and, if the series prove to be stationary or weakly persistent ($\hat{\rho} \leq 0.25$) apply the test on levels, but if even moderately persistent test on first differences.

The prevalence of revolutions

To illustrate our method we applied it to several real data sets. The first concerns a familiar subject: the spread and retreat of democracy across the globe in the course of the 20th century. In 1991 the political scientist Samuel Huntington identified three great global

"waves" of democratization (Huntington 1991). The first wave began around 1820; the second is associated with post-War War II de-colonization and the third began in 1974 and is associated with the collapse of European and Latin American dictatorships, the Iron Curtain in 1989, and the spread of democracy in Africa. Huntington evidently based his argument on a simple count of "democracies" without either defining what he meant by the term or presenting any data. Here, using much better data, we ask whether our method can identify the second and third of his waves.

To do this we use the V-Dem data set. This data set, the work of many scholars, rates the degree to which the world's nation states were democratic over the course of the 20^{th} century by means of a large number of ordinal variables that capture, in fine detail, the political structure of a given state in a given year (Coppedge et al. 2016). V-Dem provides indices where these variables have been aggregated to five higher-level quantitative variables that capture the degree to which a state exhibits: (i) freedom of expression; (ii) freedom of association; (iii) clean elections; (iv) an elected executive and, (v) universal suffrage (see Materials and Methods for details). Figure 4A (top) shows the yearly means of these variables averaged over the states extant in a given year (≤ 174). Consistent with previous V-Dem studies (Lindberg et al. 2014; Lührmann et al. 2018), it shows that global democracy has increased over the course of the 20th century but that the rate at which it has done so has not been constant. We first estimated the persistence, $\bar{\rho}$, of the series and, finding that it was ≥ 0.25 , took the first difference (S.I. Table 1). Our index, R, shows that the relative rate of change was elevated in the 1940s, early 1960s and between 1974–1999 (Figure 4B). We then carried out 3.192 significance tests over all k of which 208 were significant ($\alpha = 0.05/2$), many more than the 78 expected by chance alone, suggesting that the series contains at least one real revolution. The years in which the rate of change is significantly higher than the background rate fall into four nearly contiguous groups: 1944–1949, 1962, 1975–1985 and 1989–1996 which we then identify as distinct "revolutions" (S.I. Table 2).

Even when differenced, the entire series proved to be more persistent than desirable if we wish to avoid a high rate of Type I errors ($\bar{\rho} = 0.437$), but visual examination of the data suggested that, outside of the inferred revolutions, the series was close to stationary. To test this idea we estimated the persistence of periods before, between and after our inferred revolutions, and found that they were indeed acceptably non-persistent ($\bar{\rho} = 0.255$). We also took the second differences of the entire series, which made it overall stationary ($\bar{\rho} = -0.275$), and even so found revolutions in 1947–1948 and 1990–1992, albeit reduced in size. Thus, we are confident that the revolutions we identified are not due to the general persistence of the series.

These revolutions are very consistent with Huntington's "waves", if we allow that his



Figure 4: Cultural Revolutions. Top row of each series: Trends of individual variables. Frequency variables are shown as stacked plots; unbounded variables as lines normalised to the first time point. Middle row of each series: The rate of change index, R_i which is the sum of the F_i^k values for any time point *i* over all *k*, relative to the sum of the mean F_i^k values over all time points. Bottom row of each series: Identifying revolutions by Foote Novelty. Each cell represents the F_i^k estimate for a given half-width, *k* and time point; the colour of the cell gives the relative F_i^k value, light grey being low and dark grey being high. Note that this colour scale is only comparable within any given plot. Statistically significant ($\alpha = 0.05/2$) revolutionary periods are overlain in red, conservative periods are blue. A. global democracy (s = 5); B. pop music: Billboard Hot 100, USA (s = 100); C. Newborn girl's names, USA (s = 1423); D car models (s = 15); E. BMJ articles (s = 73); F. Old Bailey criminal trial transcripts (s = 260 — the 25% most frequent features); B. English, Irish and American novels (s = 118 — the 25% most frequent features); H crime rates per hundred thousand, UK (s = 9).

"third wave" is composed of two distinct sub-waves (c.f., Kurzman 1998; McFaul 2002; Way 2005). Interestingly, the 1962 revolution — by far the most weakly supported of the four — is an anti-democratic one caused by military coups in Indonesia, Pakistan, Greece, Nigeria, Turkey, and many Latin American countries. Huntington identified this phenomenon too and labelled it a "reverse wave" as have previous V-Dem studies (Mechkova et al. 2017). But we can add some detail to this picture. Analysis of the contributions of individual variables shows that, where the revolution of the late 1940s was due to changes in political structures, the 1977–1984 and 1989–1996 revolutions were due to an increase of personal liberty (S.I. Table 3). Revolutions, unsurprisingly, differ in their natures and causes.

Thus our method can identify times of rapid political change of the sort that political scientists and historians have spotted using less formal methods. We now turn to another familiar phenomenon: American pop music. Pop music is also said to undergo revolutionary change as new genres rise and fall, but unlike the spread of democracy there is little consensus as to when those revolutions occurred and what, exactly, changed in them (Frith 1988; Tschmuck 2006). We have previously studied the evolution of the US Billboard Hot 100, 1960–2010 (Mauch et al. 2015). In that study we assayed 17,094 songs for 16 harmonic and timbral features and, using an earlier version of our method, claimed the existence of three revolutions: in the mid-1960s, early 1980s, and late 1980s–early 1990s. We re-analysed these data using our improved testing procedure and, finding that the series is highly persistent, took the first differences. We find that $R_t > 1$ during 1967–1969, 1971, 1978, 1982–1983, 1986–1989, 1994–1995, 1998-2000, and 2005. We carried out 552 tests over all k of which 20 show a significantly elevated rate of change, more than the 14 expected by chance alone $(\alpha = 0.05/2)$; these fall into three revolutions: 1967–1969, 1982–1983, 1986–1988. These are very close to the revolutions that we previously identified and that are due, respectively, to the rise of rock-related chords and timbres (aggressive percussion) in the 1960s, the revival of guitar-heavy rock and the arrival of drum-machine percussion in the early 1980s and, in the late 1980s, the rise of hip hop at the expense of rock and pop-related timbres (S.I. Table 3). Note that since here we used differenced data, rather than levels, these are the boundaries of revolutions and not, as previously, their entire span. This accounts for the small discrepancy of dates between this analysis and the earlier one.

Besides these data sets we also applied our test to six others: the common names given to newborn girls in the USA, 1945-2010; the car models sold in the USA, 1950–2010; the articles published in the *British Medical Journal*, 1960–2008; transcripts of criminal trials at London's Old Bailey court, 1800–1900; American, Irish, and English novels published between 1840 and 1890; and a data set on the crimes committed in England and Wales 1900–2000 (See Materials and Methods for details). Of these series only the girls' names show strong evidence of revolutions particularly in the years 1973–1975 and 1988–1991 (Figure 4C–H; S.I. Table 3). These dates mark when a set of names — Jessica, Ashley, Lauren, Amanda, and Amber among others — become swiftly and immensely fashionable and then, about 15 years later, passé and replaced by names such as Emma, Isabella, Olivia and Hannah (S.I. Table 3, S.I. Figure 1). Of course, baby names change in frequency all the time (Lieberson 2000): it is the fact that several of them rose and fell in tandem that makes their dynamics revolutionary.

Although some F_t^k estimates are significant in most of the other series, they are sufficiently rare that, given the number of tests carried out, they may be due to chance alone. We wondered, however, whether some of our data sets might be susceptible to the "curse of highdimensionality" where distances between entities based on many features tend to equality (Aggarwal et al. 2001). For this reason we re-examined two high-dimension data sets, the novels and Old Bailey trials, both of which are based on more than 400 textual features expressed as probabilities or fractions, using only the 25% most common features. For the novels we still failed to find much evidence of revolutions. The Old Bailey trial data, however, now showed two strong revolutions 1819–1821 and 1833–1845. The first of these revolutions is the most interpretable. Around 1810 several synonym sets — "power", "vicarious authority / commission", "gradual change / conversion", and "smallness" — which track each other closely, increase sharply until about 1820 after which they decline; at the same time "relations to kin/consanguinity" shows an reciprocal pattern (S.I. Table 3, S.I. Figure 2). These changes may be related to a serious crisis in youth gangs at the time, one that led to a new definition of "juvenile" crime (Shore 2015). The 1834–1845 revolution appears to be more complex but may be partly related to the 1834 change in the court's jurisdiction when it became England's central criminal court (S.I. Table 3, S.I. Figure 2). The second of these periods also shows an unusual feature: although clearly a revolution when viewed at half-widths between 30 and 45, at half-widths larger than 60 it appears to be significantly conservative. This apparently paradoxical result, however, simply shows that events that are revolutionary at one scale, need not be at others.

Revolutions aside, the evolution of R_t shows interesting patterns such as a general decrease in the rate of automobile evolution 1950–2000 and a steady increase in the rate of change of crimes in England and Wales between 1960 and the 1990s. The latter is the increase of crime rates — and, for decades, their accelerating rate of increase — that occurred in Western democracies after 1960, part of what Francis Fukuyama (1991) called "The Great Disruption". We were initially surprised that we failed to identify this enormous social change as a revolution, but upon reflection it is clear that we did not since the rate increase occurred quite gradually over the course of decades. The only significant values in this series are in the late 1990s, coincident with the sharp decline in crime rates at the time. We note that all these series were persistent and so we differenced prior to analysis (S.I. Table 1); had we analysed levels instead we would have gained power and detected more revolutions but only at the cost of an increased risk of false positives.

These examples show that our method can be applied to quite different data sets: some are count data (e.g., baby names) while others are continuous traits (e.g., measure of democracy); some aggregate many individual entities that exist only in a single time interval (e.g., pop songs) while others track the evolution of a collection of entities over time (e.g., the democratic qualities of nations): all it requires is that we can estimate a distance in featurespace between intervals in a time series. Using it we have convincingly identified revolutions — some well known, others not — in several data sets, but not in all of them. This is as expected. After all, revolutions are, by definition, rare.

Conclusion

We began this paper by defining a revolution as a period of time in which the multivariate rate of change is demonstrably higher than at other times. This is most likely to occur when several variables show simultaneous increases in the rate of change. Thus our definition captures the classical idea of a revolution as a rapid, correlated, change in many properties of a system. The magnitude of change in a revolution — what we have called its strength — may be large or small in absolute terms, what matters is its size relative to the variance of change across the entire series. The period over which it occurs — what we have called its span — may be short or long.

A revolution cannot, however, span an entire time series. This is true even when all variables are changing constantly. To see this consider a collection of variables changing as directed random walks. Since each variable diverges from its original value linearly over time, its rate of change at any time, hence D over any interval, will be, within the limits of stochastic variation, constant as will F_t^k . Thus, viewed retrospectively, although there can be perpetually high rates of change, there are no perpetual revolutions. We can, however, find ourselves perpetually embroiled in revolution. When evolution is super-linear — we are thinking here of patterns such as that expressed in Moore's law of the evolution of semiconductor density (Moore 1965) — the rate of change, D over any interval, and F_t^k , all increase monotonically. In such a series a revolution will shift as the series grows so that it always defines the cutting edge. Thus there is a sense — though not, perhaps, Trotsky's (1931) — in which permanent revolutions can, and probably do, exist.

We have focused on identifying revolutions simply because times of great change capture the imagination and are invariably the subject of scholarly debate. But significance levels are, of course, arbitrary and the number of revolutions identified will change as they do. They may even be dispensed with. In their absence F_t^k , and its summary index, R_t , provides a simple way of measuring, and visualising, local variation in rates of change. We note that evolutionary biologists commonly compare rates of evolution using measures such as the darwin and the haldane. Although both can be applied to any kind of time-series data, both are univariate and generally estimated over an entire series (Lambert et al. 2017), and so not well suited to estimating temporal variation in rates of multivariate evolution.

In all our data sets, all variables had non-zero values. However, it is possible to imagine revolutions in which some variables become irrelevant even as others arise. To give a concrete example, consider car design. Over fifty years of car evolution we detected much change, but no revolutions. Now, however, electric cars are upon us. Some of their features are much like those of their fossil-fuelled ancestors (e.g., door number), but some (e.g., cylinder number, gear number) are not applicable, others can still be measured but are radically different (e.g., the relationship between maximum torque and RPM), while yet others are altogether new (e.g., power train battery capacity). Such changes in the salience of variables can be handled by our method and, if they have a sufficiently swift and strong effect on the multivariate distribution, will appear as a revolution. However, the revolution they will surely bring about seems to be of a different kind than any involving merely quantitative changes, however rapid, in mean horsepower or chassis length. The fundamental distinction is between revolutions that entail changes in the relationships among variables or, more formally, their variance-covariance structure, and those that do not. We think of the former as "structural" revolutions (c.f. Snodgrass 1980) and the latter "non-structural", note that they are subsets of the revolutions that our method detects, but leave the problem of telling them apart for future research.

Our method can be applied to identifying dramatic changes in any multivariate time series data of sufficient length and quality. In biology it might be applied to the study of gene expression profiles, the evolution of gene frequencies or morphology (e.g., Tu et al. 2005; Bergland et al. 2014; Hunt et al. 2015). But the idea of revolution has its origin in historiography and so we have focused on political, social and cultural phenomena. As large data sets capturing their evolution become available, it is increasingly possible to use statistical inference to test historical hypotheses (e.g., Michel et al. 2011; Hughes et al. 2012; Rodriguez Zivica et al. 2013; Perc 2013; Klingenstein et al. 2014; Rule et al. 2015; Bearman 2015). It seems to us that this methodological change is so profound that future historians of history may even call it a revolution.

Materials & Methods

Foote Novelty

Foote Novelty estimation, significance testing, and other procedures were implemented in R; code is available from the authors by request.

Foote Novelty performance

We assessed the ability of our algorithm to detect revolutions in artificial data series where the revolutionary span and strength were known. Since most series encountered will have a degree of persistence (ρ) which lies somewhere between a white noise process ($\rho = 0$) and a random walk ($\rho = 1$), we use the following time series process to interpolate between these extrema,

$$X_{i,t} = \begin{cases} \rho X_{i,t-1} + \epsilon_t, \\ \text{if } t < t_{start} \\ r_i + r_i(t - t_{start} - 1)(1 - \rho) + \rho X_{i,t-1} + \epsilon_t \\ \text{if } t \in t_{rev} \\ r_i(t_{end} - t_{start})(1 - \rho) + \rho X_{i,t-1} + \epsilon_t, \\ \text{if } t > t_{end} \end{cases}$$

where t_{start} and t_{end} are the periods when the revolution begins and ends, respectively; t_{rev} consists of all time points throughout the course of the revolution; $\epsilon_t \sim \mathcal{N}(0, \sigma)$. We estimated the probability of Type I (rate of false positives) and Type II (rate of false negatives) errors that resulted from applying the FN algorithm to time series generated by process (1). We varied the number of variables in our simulations, between 0 and 200 and, for each series, we allowed a separate revolution displacement (although for all variables within a replicate, the revolutions always began and ended at the same times). Specifically, the revolution size for variable *i* is drawn from a normal distribution, $r_i \sim \mathcal{N}(0, s)$, where *s* represents the 'revolution strength'. We evaluated the performance of our revolution detection algorithm across a range of parameter values, including the Foote Novelty kernel half-width, the number of variables in our dataset, and the revolution strength and duration. Each simulation was run for 100 time points, and in all cases the revolutions began at the 40th time point. We used the following parameter combinations: persistence (ρ): 0, 0.2, 0.4, 0.6, 0.8, 1.0; revolution strength (*s*): 0.0, 0.2, 0.4, 0.6, 0.8, 1.0, 1.2, 1.4, 1.6, 1.8, 2.0; revolution length (l = % of series length): 2, 6, 10, 14, 18; variable number (*n*): 10, 30, 50, 70, 90, 110, 130, 150, 170, 190. For these 3850 series we estimated F_i^k at kernel half-widths (k): 1, 3, 5, 7, 9. To estimate the rate of Type I and Type II errors we used 20 replicates at each parameter set. In each case, we applied our algorithm to the collection of variables, and recorded whether we detected at least one revolutionary period outside the true period (a false positive), or failed to detect a revolution anywhere in the period when it occurred (a false negative). The results of these simulations are given in Figure 3.

Estimating persistence, ρ_i

To estimate the persistence of the real time series, we follow (Lambert et al. 2017), which amounts to estimating a hierarchical Bayesian model of the form,

$$X_{i,t} \sim \mathcal{N}(\rho_i X_{i,t-1}, \sigma_i),$$

$$\rho_i \sim \mathcal{N}(\rho', \sigma'), \tag{1}$$

where the population-level parameters are assigned priors: $\rho' \sim \mathcal{N}(0,1)$ and $\sigma' \sim \text{half-}\mathcal{N}(0,1)$.

Data

The pop song, clinical article, novel, car model and Old Bailey data sets consist of many individual artefacts, each of which is represented only once in the data set, at its date of first appearance. The properties of these artefacts are measured, and the analysis is based on the aggregate properties of the population. The democracy, crimes, baby names data sets consist of categories present throughout the series. The analysis is based on the relative values, incidence or frequencies over time. Most of the data that we used have been published elsewhere, so we only sketch their provenance. When computing pairwise distances among years either Euclidean distance or Kullback-Leibler distance (a symmetrized version of KL divergence) was used as appropriate; variables were scaled to have a mean = 0 and standard deviation = 1 when appropriate.

Pop music. 17,094 unique songs comprising about 80% of the population of the US Billboard Hot 100 1960-210. The traits are 100 harmonic and timbral topics in each. (Mauch et al. 2015; Lambert et al. 2017). Clinical articles. 170,577 clinical articles from *British Medical Journal* between 1960 and 2008. The traits are 73 Topics (Mauch et al. 2015; Lambert et al. 2017). Novels. 2,203 novels collected by the Stanford Literary Lab in 2015. The traits are 471 Topics (Jockers 2013; Lambert et al. 2017). Cars. 2,210

car models sold in the USA 1950-2010. The traits are 16 quantitative variables describing the power train performance and car dimensions (Lambert et al. 2017). Democracy. Based on the https://www.v-dem.net/en/data/data-version-7-1/ V-Dem dataset (Version 7.1) which describes the political organization of up to 174 countries between 1900 and 2015. We focused on five aggregate variables that collectively describe the degree of "polyarchy" or, more colloquially, democracy. These variables measure freedom of expression (e_v2x_freexp_thick_5C), freedom of association (e_v2x_frassoc_thick_5C), share of population with suffrage $e_v2x_suffr_5C$, clean elections ($e_v2xel_frefair_5C$), and elected executive (e_v2x_accex_5C) which, in turn, are aggregates of other subordinate variables (Coppedge et al. 2016; Pemstein et al. 2017). Crime. Based on https://data.gov.uk/dataset/recordedcrime-data-1898-2001-02/resource/b5b1c3fe-338e-472e-b844-75108c57436c crime statistics from the UK Home Office website which records 154,300,472 crimes committed in England and Wales 1898–2002. We truncated these to 1900–2000 and normalised them by population size to give crime rates per 100,000 residents. The Home Office notes that the classification of crimes has varied over time. For this reasons we used the nine summary categories: "homicide + manslaughter", "sexual", "robbery", "violence", "burglary", "theft", "fraud", "other" which are largely immune to these changes. In any event, the increase of crime rates after the late 1960s is a well known phenomenon not due to changes in the law or reporting (Fukuyama 1991). Names. The most common 1,423 names given to newborn girls in the USA 1945–2015, subsetted from the https://www.ssa.gov/oact/babynames/limits.html data collected by the US Social Security Administration. Only names present throughout the period were used. Old Bailey trial records. 112,485 trial records recorded between 1760 and 1913 at the Central Criminal Court, or Old Bailey, in London, a period during which trial reports were at their most comprehensive. The traits are the frequencies of 1,040 synonym sets based on the 20 million (semantic) words of testimony (Klingenstein et al. 2014).

Authors' contributions

A.M.L., B.L., and M.M. designed the study, contributed new reagents / analytic tools, collected data, carried out analysis and wrote the paper. M.P. contributed new reagents/-analytic tools; P.L., S.L. and S.A. contributed data. All authors gave final approval for publication.

Competing interests

All authors declare no competing interests

Funding

B.L. was supported by EPSRC grant code: EP/F500394/1. V-Dem data collection was supported by European Research Council, Grant 724191; Riksbankens Jubileumsfond, Grant M13-0559:1; Knut and Alice Wallenberg Foundation to Wallenberg Academy Fellow, Grant 2013.0166; as well as by internal grants from the Vice- Chancellor's office, the Dean of the College of Social Sciences, and the Department of Political Science at University of Gothenburg and the Marianne and Marcus Wallenbergs Foundation Grant 2017.0049.

Acknowledgements

We thank Simon deDeo, Tim Hitchcock and Matthew Jockers for supplying data. Arnold Heumakers, Tim Hitchcock, James Rosindell and Ted Underwood for comments on the manuscript.

References

- Aggarwal, C. C., A. Hinneburg, and D. A. Keim (2001). "On the Surprising Behavior of Distance Metrics in High Dimensional Space". In: Springer Berlin Heidelberg, pp. 420–434.
- Bearman, P. (2015). "Big Data and historical social science". In: *Big Data & Society* 2.2, pp. 1–5.
- Bergland, A., E. Behrman, K. O'Brien, P. Schmidt, and D. Petrov (2014). "Genomic Evidence of Rapid and Stable Adaptive Oscillations over Seasonal Time Scales in Drosophila". In: *PLoS Genetics* 10, e1004775.
- Bezanson, A. (1922). "The Early Use of the Term Industrial Revolution". In: *The Quarterly Journal of Economics* 36.343-349.
- Bowler, P. (1988). The Non-Darwinian Revolution. Reinterpreting a Historical Myth. Johns Hopkins University Press.
- Butterfield, H. (1950). The Origins of Modern Science. G. Bell and Son.
- Cohen, H. F. (1994). The Scientific Revolution: A historiographical inquiry. University of Chicago Press.

Cohen, I. B. (1986). Revolution in Science. Harvard.

- Coppedge, M., J. Gerring, S. I. Lindberg, S.-E. Skaaning, J. Teorell, D. Altman,
 M. Bernhard, M. S. Fish, A. Glynn, A. Hicken, C. H. Knutsen, K. Marquardt,
 K. McMann, P. Paxton, D. Pemstein, J. Staton, B. Zimmerman, F. Andersson,
 V. Mechkova, and F. Miri (2016). V-Dem Country-Year Dataset v6. V-Dem Institute,
 Gothenburg: Varieties of Democracy (V-Dem) Project.
- De Ruiter, R. and J. Schalk (2017). "Explaining cross-national policy diffusion in national parliaments: A longitudinal case study of plenary debates in the Dutch parliament". In: *Acta Politica* 52.2, 133–155.
- Eldredge, N. and S. J. Gould (1972). "Punctuated equilibria: An alternative to phyletic gradualism". In: *Models in Paleobiology*. Ed. by T. Schopf. New York: Freeman Cooper.
- Flink, C. M. (2017). "Rethinking Punctuated Equilibrium Theory: A Public Administration Approach to Budgetary Changes". In: *Policy Studies Journal* 45, 101–120.
- Foote, J. (2000). "Automatic audio segmentation using a measure of audio novelty". In: *IEEE International Conference on Multimedia and Expo.* Vol. 1, pp. 452–455.
- Fowler, L., T. T. Neaves, J. N. Terman, and A. G. Cosby (2017). "Cultural Penetration and Punctuated Policy Change: Explaining the Evolution of US Energy Policy". In: *Review of Policy Research* 34, 559–577.
- Frith, S. (1988). Music for Pleasure. Cambridge, UK: Polity Press.
- Fukuyama, F. (1991). The Great Disruption: human nature and the reconstitution of social order. Profile.
- Gould, S. J. (2002). *The structure of evolutionary theory*. Cambridge, MA: Harvard University Press.
- Hall, A. R. (1954). The Scientific Revolution, 1500-1800: The formation of the modern scientific attitude. Beacon Press.
- Heumakers, A. (2015). De Esthetische Revolutie. Boom.
- Himmelfarb, G. (1996). Darwin and the Darwinian Revolution. Ivan R. Dee.
- Hodge, M. (2005). "Against "Revolution" and "Evolution". In: Journal of the History of Biology 38.101-121.
- Hughes, J. M., N. J. Foti, D. C. Krakauer, and D. N. Rockmore (2012). "Quantitative patterns of stylistic influence in the evolution of literature". In: *Proceedings of the National Academy of Sciences*, USA 109.20, pp. 7682–7686.
- Hunt, G., M. J. Hopkins, and S. Lidgard (2015). "Simple versus complex models of trait evolution and stasis as a response to environmental change". In: *Proceedings of the National Academy of Sciences*, USA 112.16, 4885–4890.

Huntington, S. (1991). "Democracy's third wave". In: Journal of Democracy 2, pp. 12–34.

- Jockers, M. (2013). *Macroanalysis: Digital methods and literary history*. University of Illinois Press.
- Klingenstein, S., T. Hitchcock, and S. DeDeo (2014). "The civilizing process in London's Old Bailey". In: Proceedings of the National Academy of Sciences 111.26, pp. 9419–9424.
- Koyré, A., ed. (1957). From the Closed World to the Infinite Universe. Johns Hopkins Press.
- Kuhn, T. (1957). The Copernican Revolution: Planetary astronomy in the development of western thought. Harvard University Press.
- (1972). The Structure of Scientific Revolutions. University of Chicago Press.
- Kurzman, C. (1998). "Waves of democratization." In: Studies in Comparative International Development 33.42–64.
- Lambert, B., G. Kontonatsios, M. Mauch, T. Kokkoris, M. Jocker, S. Ananiadou, and A. Leroi (2017). "The pace of modern culture". In: *in review Proceedings of the National Academy of Sciences, USA* 00, pp. 00–00.
- Lieberson, S. (2000). A matter of taste: how names, fashions and culture change. Yale.
- Lindberg, S. I., M. Coppedge, J. Gerring, and J. Teorell (2014). "A New Way to measure Democracy". In: *Journal of Democracy* 25, pp. 159–169.
- Lonne, I. (2017). "Terrestrial slopes in northern high latitudes: A paradigm shift regarding sediment origin, composition, and dynamic evolution". In: *Geomorphology* 276, pp. 180–202.
- Lowenstein, J. and J. J. Grantham (2017). "Residual renal function: a paradigm shift". In: *Kidney International* 91.3, pp. 561–565.
- Lührmann, A., V. Mechkova, M. Olin, F. P. Casagrande, C. S. Petrarca, L. Saxer, and S. I. Lindberg (2018). *Democracy at Dusk? V-Dem Annual Report 2017*. Tech. rep. V-Dem Institute, University of Gothenburg.
- Mauch, M., R. M. MacCallum, M. Levy, and A. M. Leroi (2015). "The evolution of popular music: USA 1960–2010". In: *Royal Society Open Science* 2, pp. 150081+.
- Mcbrearty, S. and A. S. Brooks (2000). "The revolution that wasn't: a new interpretation of the origin of modern human behavior". In: *Journal of Human Evolution* 39, pp. 453–563.
- McFaul, M. (2002). "The Fourth Wave of Democracy and Dictatorship: Noncooperative Transitions in the Postcommunist World". In: *World Politics* 54.212–244.
- Mechkova, V., A. Lührmann, and S. I. Lindberg (2017). "How Much Democratic Backsliding?" In: *Journal of Democracy* 28, pp. 162–169.
- Mellars, P. and C. Stringer (1989). The Human Revolution: Behavioural and Biological Perspectives on the Origins of Modern Humans. Edinburgh University Press.

- Michel, J. B., Y. K. Shen, A. P. Aiden, A. Veres, M. K. Gray, T. G. B. Team, J. P. Picket, D. Hoiberg, D. Clancy, P. Norvig, J. Orwant, S. Pinker, M. A. Nowak, and E. L. Aiden (2011). "Quantitative analysis of culture using millions of digitized books". In: *Science* 331, pp. 176–182.
- Moore, G. E. (1965). "Cramming More Components onto Integrated Circuits". In: *Electronics* 38, pp. 114–117.
- Omranian, N., B. Mueller-Roeber, and Z. Nikoloski (2015). "Segmentation of biological multivariate time-series data". In: *Scientific Reports* 5, pp. 8937–.
- Pemstein, D., K. L. Marquardt, E. Tzelgov, Y.-t. Wang, J. Krusell, and F. Miri (2017). The V-Dem Measurement Model: Latent Variable Analysis for Cross-National and Cross-Temporal Expert-Coded Data. Tech. rep. Varieties of Democracy Institute: Working Paper No. 21.
- Pennel, M. W., L. J. Harmon, and J. C. Uyeda (2013). "Is there room for punctuated equilibrium in macroevolution?" In: *Trends in Ecology and Evolution* 29, pp. 23–30.
- Perc, M. (2013). "Self-organization of progress across the century of physics". In: Scientific Reports 3.
- Preuss, P., R. Puchstein, and H. Dette (2015). "Detection of Multiple Structural Breaks in Multivariate Time Series". In: *Journal of the American Statistical Association* 110.510, pp. 654–668.
- Raoult, D. (2017). "Double paradigm shift for the antibiotics' activity on viruses: Zika's lesson". In: *Proceedings of the National Academy of Sciences* 114.7, E1045.
- Rodriguez Zivica, P. H., F. Shifres, and G. A. Cecchic (2013). "Perceptual basis of evolving Western musical styles". In: *Proceedings of the National Academy of Sciences*, USA 110, pp. 10034–10038.
- Rule, A., J.-P. Cointet, and P. S. Bearman (2015). "Lexical shifts, substantive changes, and continuity in State of the Union discourse, 1790–2014". In: *Proceedings of the National Academy of Sciences* 112.35, pp. 10837–10844.
- Ruse, M. (1979). *The Darwinian Revolution: Science Red in Tooth and Claw.* The University of Chicago Press.
- Schama, S. (1989). Citizens: A Chronicle of The French Revolution. Viking.
- Sepkoski, D. (2012). Rereading the fossil record: the growth of paleobiology as an evolutionary discipline. University of Chicago Press.
- Sepkoski, D. and M. Ruse, eds. (2009). *The Paleobiological Revolution: Essays on the Growth of Modern Paleontology*. University of Chicago Press.
- Seward, J. B. (2017). "Paradigm Shift in Medical Data Management: Big Data and Small Data?" In: JACC: Cardiovascular Imaging,

Shapin, S. (1996). The Scientific Revolution. University of Chicago Press.

- Shore, H. (2015). London's Criminal Underworlds, c. 1720–c. 1930: A Social and Cultural History. Palgrave Macmillan.
- Snodgrass, A. (1980). Archaic Greece: The Age of Experiment. University of California Press.
- Trotsky, L. (1931). The Permanent Revolution. Progress Publishers.
- Tschmuck, P. (2006). Creativity and innovation in the music industry. Dordrecht: Springer.
- Tu, B. P., A. Kudlicki, M. Rowicka, and S. L. McKnight (2005). "Logic of the Yeast Metabolic Cycle: Temporal Compartmentalization of Cellular Processes". In: *Science* 310.5751, pp. 1152–1158.
- Way, L. (2005). "Authoritarian State Building and the Sources of Regime Competitiveness in the Fourth Wave: The Cases of Belarus, Moldova, Russia, and Ukraine." In: World Politics, 57.231–261.

Supplementary Information

data set	$ar{ ho}_u$	d.o.	$ar{ ho}_d$
global democracy (whole)	0.998	1	0.437
global democracy (non. rev.)	0.913	1	0.255
global democracy (whole)	0.998	2	-0.275
US pop music	0.875	1	-0.15
US girls' names	0.961	1	0.070
US cars	0.629	1	-0.528
BMJ	0.814	1	-0.323
Old Bailey trials	0.573	1	-0.432
English, Irish & US novels	0.396		
English & Welsh crime	1.055	1	0.235

S.I. Table 1: Estimates of overall persistence on undifferenced, $\bar{\rho}_u$, and differenced $\bar{\rho}_d$, series. To avoid a rate of Type I errors greater than the set significance value, $\alpha = 0.05$, we require series to be have low persistence, $\rho \leq 0.25$ (Main Text Figure 3), and difference until we achieve that. The *d.o.* column gives the degree of differencing. Since the global democracy series had a $\rho > 0.25$ even after first differencing we examined it further by, first, looking at the persistence of non-revolutionary ("non-rev") periods and, second, by taking the second differences of the whole series. In both cases we found near-stationarity showing that the inferred revolutions are not due to the series being generally persistent. See Materials & Methods for details.

data set	revolution	years	k
global democracy	1	1944 - 1948	2 - 20
	2	1962	16
	3	1977 - 1984	14 - 20
	4	1989 - 1996	1 - 8
US pop music	1	1968	4-6
	2	1982 - 1983	1 - 3
	3	1986 - 1988	3 - 14
US girls' names	1	1974 - 1975	5 - 17
	2	1988 - 1991	8 - 20
Old Bailey trials	1	1819 - 1821	1 - 45
	2	1833 - 1845	32 - 46

S.I. Table 2: Summary of revolutions detected with FN novelty. Significant FN values were found in most other series, but only in these series were the number of significant tests greater than those expected by chance alone. Note the Old Bailey results are based on the most frequent 25% of traits alone; considering all the traits shows no evidence of revolutions.

data set	revolution	years	trait	direction
global democracy	1	1944 - 1948	free and fair elections	1
	1		share of population with suffrage	1
	1		elected executive	1
	3	1977 - 1984	freedom of expression	↑
	3		freedom of association	↑
	4	1989 - 1996	freedom of association	↑
	4		freedom of expression	↑
US pop music	1	1968	h3 – minor 7th chords	<u>^</u>
			h4 – standard diatonic chords	↑
			t1 – drums, aggressive, percussive	↑
	2	1982 - 1983	t1 – drums, aggressive, percussive	^
			t5 - guitar, loud, energetic	↑
	3	1986 - 1988	t1 – drums, aggressive, percussive	. .
	ũ.		t_3 – energetic, speech, bright	\uparrow
			t7 - /oh/, rounded, mellow	
US girls' names	1	1974-1975	Jessica	<u>↓</u> ↑
	Ŧ	1011 1010	Amanda	↑
			Sarah	↑
			Amber	↑
			Ashley	↑
			Lauren	↑
			Michelle	↑
			Amy	↑
			Iamio	↑
			Mary	
			Megan	↓ ↑
	2	1088-1001	Jaccica	
	2	1500 1551	Fmma	↓ ↑
			Ashlov	
			Amanda	↓
			Allovia	↓
			Icabella	 ★
				 ★
			Ulivia	 ★
			Mami	
Old Deilers triels	1	1010 1001	Mary	<u>↓</u> *
Old Balley trials	1	1819-1821	ssii - relations of kindred, consanguinity	
			ss_{157} – power	↓ ↓
			ss755 - vicarious authority / commission	↓ ↓
			ss144 - gradual change to something different / conversion	+
	0	1099 1045	$ss_{32} - sinaliness$	\downarrow
	2	1833 -1845	$ss_{194} - parter$	\downarrow
			ssoo – nonincrease/decrease	↓ ★
			ssoup – irresolution	
			$ss_{32} - smallness$	↓
			ssybb – jurisdiction / executive	\uparrow
			ss541 – learner	$\downarrow \uparrow$
			ss570 – perspicuity	\downarrow
			ss995 – churchdom	\uparrow

S.I. Table 3: Variables that contribute most to revolutions ordered in terms of importance from high to low . US girls' names: Only names that, when removed, decreased the number of significant F_i^k values by $\geq 25\%$ in a given revolution are listed. Old Bailey trials: Only variables that, when removed, decreased the number of statistically significant F_i^k values by $\geq 35\%$ in a given revolution are listed.



S.I. Figure 5: US girls' names: Only names that, when removed, decreased the number of significant F_i^k values by $\geq 25\%$ in a given revolution are listed.



S.I. Figure 6: Old Bailey trials: variables that contributed most to the 1819-1821 and 1833-1844 revolutions. Only variables that, when removed, decreased the number of statistically significant F_i^k values by $\geq 35\%$ in a given revolution are listed. The 1819–1821 revolution is largely driven by variables related to changes in the treatment of juvenile crime; the 1834-1844 revolution is partly driven by changes in court jurisdiction.